

Cusps Enable Line Attractors for Neural Computation

Zhuocheng Xiao^{1,2}, Jiwei Zhang³, Andrew T. Sornborger^{4,*} and Louis Tao^{1,5,*}

¹*Center for Bioinformatics, National Laboratory of Protein Engineering and Plant Genetic Engineering, College of Life Sciences, Peking University, Beijing, 100871, China*

²*Department of Mathematics, University of Arizona, Tucson, AZ 85721, USA*

³*Beijing Computational Science Research Center, Beijing, 100193, China*

⁴*Department of Mathematics, University of California, Davis, CA 95616, USA and*

⁵*Center for Quantitative Biology, Peking University, Beijing, 100871, China*

(Dated: February 22, 2017)

Line attractors in neural networks have been suggested to be the basis of many brain functions, such as working memory, oculomotor control, head direction, locomotion, and sensory processing. In recent work, we incorporated pulse gating into feedforward neural networks and showed that the transmission of graded information can be viewed as a line attractor in the firing rate of transiently synchronous populations. While this was revealed in an analysis using rate models, graded information transfer persisted in spiking neural networks and was robust to intrinsic and extrinsic noise. Here, using a Fokker-Planck approach, we show that the gradedness of rate amplitude information transmission in pulse-gated networks is associated with the existence of a cusp catastrophe, and that the slow dynamics near the fold of the cusp underlies the robustness of the line attractor. Understanding the dynamical aspects of this cusp catastrophe allows us to show how line attractors can persist in biologically realistic neuronal networks and how the interplay of pulse gating, synaptic coupling and neuronal stochasticity can be used to enable attracting one-dimensional manifolds and thus, dynamically control graded information processing.

PACS numbers: 87.18.Sn,87.19.lj,87.19.lm,87.19.lq,87.19.ls,05.10.Gg

Attractor neural networks have been appealed to by theoretical neuroscientists to explain working memory [1], oculomotor control [2], head direction [3], locomotion [4], sensory processing [5] and many other experimentally observed brain functions. Some of these neuronal functions have been rationalized using the persistent activity of an attractor (e.g., working memory), while other functions have been modeled by the gradedness (i.e., order preserving property [6]) of an attractor (e.g., oculomotor control and sensory processing). Because of these dynamical features, attractor networks have been particularly useful in modeling the encoding of continuous external stimuli [7] or the maintenance of internal neuronal representations [6]. Furthermore, it has been shown that attractor network dynamics have the capacity to perform optimal computations [8] and thus, can implement Bayesian inference [9, 10].

Pulse-gating is a mechanism capable of transferring packet-based spiking activity from one neural population to another [11, 12]. For instance, the synfire-gated synfire chain (SGSC) [12] is a mechanism consisting of two sets of neural populations, one representing rate-coded (graded) information and another that gates the flow of that information. This separation into information-carrying and information-control populations provides a basis for understanding how some brain areas can be responsible for the control of information transfer [13], such as the mediodorsal nucleus of the thalamus [14], and others for information processing, such as the somatosensory areas [15].

Using a mean-field, rate model analysis, we showed that graded firing rate amplitudes may be exactly propagated from one neuronal population to another using a pulse-gating mechanism [11, 12]. Consequently, the propagation of nearly synchronous, transiently spiking populations can be viewed

macroscopically as a line attractor in the amplitude of spiking activity packets. The graded transfer of firing amplitudes was subsequently observed in numerical simulations of spiking neuronal networks and was robust to intrinsic, external and parametric noise.

In order to understand pulse-gated, graded transfer in spiking models, we used a Fokker-Planck approach to study the state space of this line attractor [16]. We demonstrated the dynamical convergence of membrane potential density to a one-dimensional manifold. We were able to show that, in the mean, pulse-gating in feedforward networks gives rise to nearly time-translationally-invariant spiking probabilities that are propagated from layer to layer in a feedforward network. Numerical simulations of our Fokker-Planck system revealed a saddle-node bifurcation where the unstable fixed point is strongly attracting and weakly unstable. Thus, on the proper time-scale, the unstable manifold of the saddle can be viewed as an attracting one-dimensional manifold.

Here, we show that this approximate line attractor is robust and generic. By using a reduced model of our Fokker-Planck system, we show that the propagation of firing amplitudes can be mapped to a cusp catastrophe, and that a bundle of one-dimensional attracting manifolds exist in the region surrounding the fold of the cusp. Our results reveal how the coordination of pulse-gating, synaptic coupling, and membrane potential dynamics enables these line attractors in feedforward networks and demonstrate the robustness of graded information propagation when pulse-gating is incorporated.

We studied a feedforward network of $j = 1, \dots, M$ populations of $i = 1, \dots, N$ excitatory, current-based, integrate-and-fire (I&F) neurons whose membrane potential, V , and

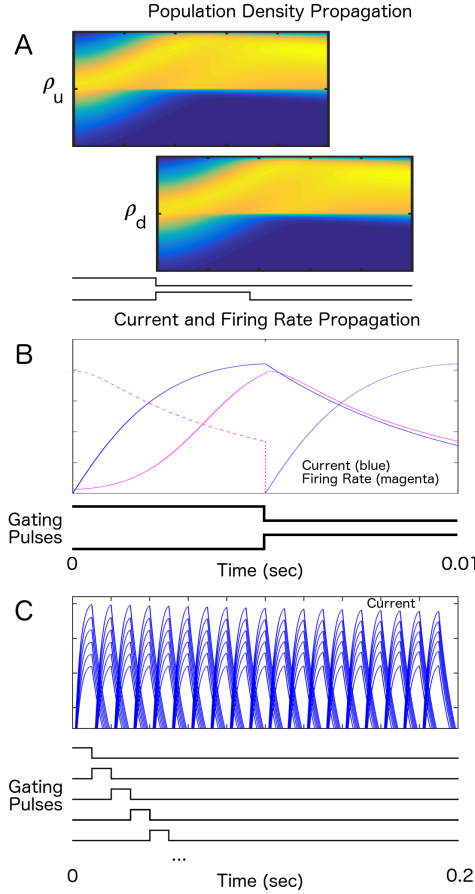


FIG. 1. Graded, pulse-gated propagation in a neural Fokker-Planck model. A) Graded population density transfer. Population densities in upstream, $\rho_u(V, t)$, and downstream, $\rho_d(V, t)$, layers. These are two layers from a 19 layer feedforward network with $S = 2.83$, $I_{gate} = 14.2$, and $\sigma_0 = 4.5$. B) Synaptic current, \bar{I}_j^{ff} (blue), and firing rate, $m_j(t)$ (magenta), in upstream and downstream layers. Below the plot, we show the timing of the gating pulses. C) Graded firing rates showing faithful, pulse-gated propagation across 19 layers. Gating pulses for the first few layers are shown below the plot.

synaptic current are described by

$$\frac{d}{dt}V_{i,j} = -g_L(V_{i,j} - V_R) + I_j^g + I_{i,j}^{ff} \quad (1a)$$

$$\tau \frac{d}{dt}I_{i,j}^{ff} = -I_{i,j}^{ff} + \begin{cases} \frac{S}{pN} \sum_k p_{jk} \sum_l \delta(t - t_{j-1,k}^l), & j > 1 \\ A \delta(t), & j = 1 \end{cases} \quad (1b)$$

where V_R is the reset voltage, τ is the synaptic timescale, S is the synaptic coupling strength, p_{jk} is a Bernoulli distributed random variable and $p = \langle p_{jk} \rangle$ is the mean synaptic coupling probability. The l 'th spike time of the k 'th neuron in layer $j - 1$ is determined by $V_{i,j}(t_{j-1,k}^l) = V_{Th}$, i.e. when the neuron reaches threshold (after which $V_{i,j}$ is immediately reset to V_R). The gating current, I_j^g , is a white noise process with a square pulse envelope, $\theta(t - (j - 1)T) - \theta(t - jT)$, where θ is a Heaviside theta function and T is the pulse length [11] of pulse height \bar{I}^g and variance σ_0^2 . Note that an exponentially decaying current is injected in population 1 providing synchronized activity that will subsequently propagate downstream through populations $j = 2, \dots, M$.

Feedforward networks of this type admit discrete-time-translationally invariant solutions describing graded packet transfer [16]. Such solutions exist for gating pulses that are either temporally sequential or overlapping [12].

To understand this phenomenon, we used a Fokker-Planck analysis. Assuming the spike trains in Eq. (1b) to be Poisson distributed, the collective behavior of this feedforward network may be described by the (Fokker-Planck) equations

$$\frac{\partial}{\partial t} \rho_j(V, t) = -\frac{\partial}{\partial V} J_j(V, t) \quad (2a)$$

$$\tau \frac{d}{dt} \bar{I}_j^{ff} = -\bar{I}_j^{ff} + \begin{cases} S m_{j-1}, & j > 1 \\ A \delta(t), & j = 1 \end{cases} \quad (2b)$$

These equations describe the evolution of the probability density function, $\rho_j(V, t)$, in terms of the probability density flux, $J_j(V, t)$, the mean feedforward synaptic current, \bar{I}_j^{ff} , and the population firing rate, m_j . For each layer, j , the probability density function gives the probability of finding a neuron with membrane potential $V \in (-\infty, V_{Th}]$ at time t .

The probability density flux is given by

$$J_j(V, t) = \left([-g_L(V - V_R) + \bar{I}^g + \bar{I}_j^{ff}] - D_j \frac{\partial}{\partial V} \right) \rho_j(V, t).$$

The effective diffusivity is

$$D_j = \sigma_0^2 + \frac{1}{2} \frac{S^2}{pN} m_{j-1}(t). \quad (3)$$

(Below, we take $N \rightarrow \infty$ and neglect the second term on the right in (3).) The population firing rate is the flux of the probability density function at threshold,

$$m_j(t) = J_j(V_{Th}, t). \quad (4)$$

The boundary conditions for the Fokker-Planck equations are $J_j(V_R^+, t) = J_j(V_{Th}, t) + J_j(V_R^-, t)$, $\rho_j(V_R^+, t) = \rho_j(V_{Th}, t) + \rho_j(V_R^-, t)$, and $\rho_j(V = -\infty, t) = 0$.

In Fig. 1, we show that population density (Fig. 1A), and graded current and firing rate (Fig. 1B) may be propagated between layers in the Fokker-Planck model. In Fig. 1C, we demonstrate stable graded propagation across many layers.

In the graded transfer regime, the probability density may be described by a few of its low-lying moments, *viz.* $M_j^1(t) = \int_{-\infty}^{V_{Th}} V \rho_j(V, t) dV$ and $M_j^2(t) = \int_{-\infty}^{V_{Th}} V^2 \rho_j(V, t) dV$. Changes in the moments are related to changes in the shape of the density. To identify time-translationally invariant solutions admitting graded propagation, we search for parameters, S , A , etc. such that for a given T , $\bar{I}_j^{ff}(t + T) \approx \bar{I}_{j-1}^{ff}(t)$.

In Fig. 2, we plot the map $(m_j(t + jT), M_j^1(t + jT), M_j^2(t + jT))$ for successive values of j (i.e. mapping feedforward propagation from layer to layer). As may be seen in Fig. 2, for parameters admitting graded information propagation, initial conditions rapidly approach a one-dimensional manifold. On the manifold itself, there are three fixed points, two stable fixed points at the extremes and one unstable saddle in between. For these parameters, the dynamics along the unstable direction of the saddle are slow. On the other hand, the dynamics along the stable direction of the saddle

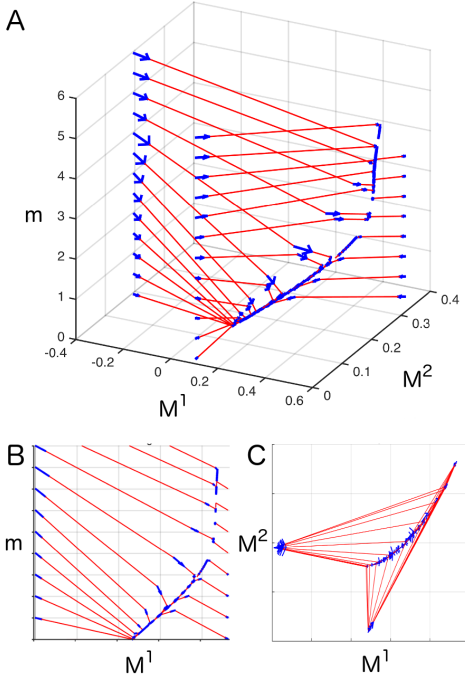


FIG. 2. The slow manifold of graded propagation. A) Trajectories, $(m_j(t+jT), M_j^1(t+jT), M_j^2(t+jT))$ for successive values of j , showing the rapid approach of the initial distribution to a one-dimensional manifold and the slow evolution along the manifold representing the slow decay of the waveform. B) The trajectories in A) projected onto the $M^1 - m$ plane. C) The trajectories in A) projected onto the $M^1 - M^2$ plane. Parameters are as in Fig. 1.

are fast. Thus, after an initial transient, the amplitude and waveform of each packet remain relatively unchanged from transfer to transfer. Furthermore, the relative ordering of the amplitudes is retained even as the amplitude slowly changes.

Fixed points of the map $(m_j(t+jT), M_j^1(t+jT), M_j^2(t+jT))$ are shown in Fig. 3, showing a fold bifurcation (along the unstable manifold of the saddle) as a function of S .

As S moves away from the region of graded transfer, the dynamics along the unstable manifold become fast. Thus, all amplitudes rapidly approach the fixed points at the extremes. This leads to binary transfer, where pulse gating can only transfer low and high amplitudes. As S increases, the location of the high fixed point increases giving more distinction between low and high fixed points.

This behavior may be understood with a model in which the initial distribution is assumed to be Gaussian, $\rho_0(V, t) = (1/P) \exp(-(V - \mu(t))^2/2\sigma^2)$, with width σ and mean $\mu(t)$, where $P = \int_{-\infty}^{V_{Th}} \rho_0(V, 0)$ is a normalization factor accounting for the truncation of the Gaussian at threshold, V_{Th} . As the gating current turns on, the distribution is advected toward the voltage threshold, V_{Th} , and the population begins to fire. Since the timescale of the pulse is fast, neurons only have enough time to fire once (approximately). Thus, we neglect the re-emergent population at V_R . With this approximation, Eq. (2a) gives rise to $\dot{\mu} = -g_L(\mu - V_R) + \bar{I}^g + \bar{I}_u^{\text{ff}}$, where $\sigma^2 = \sigma_0^2/g_L$. With upstream current $\bar{I}_u^{\text{ff}} = Ae^{-t/\tau}$. Setting




FIG. 3. Bifurcation diagram generated using numerically integrated Fokker-Planck equations (2), for $I = 14.2$. The saddle-node bifurcation occurs at approximately $S = 2.82$. For values close to $S = 2.82$ the trajectory evolves slowly along the one-dimensional manifold due to a fold bifurcation ghost, allowing graded propagation across many layers before significant change in the waveform.

$V_{Th} = 1$, this integrates to

$$\mu(t) = \mu_0 e^{-g_L t} + \frac{\bar{I}^g}{g_L} (1 - e^{-g_L t}) + \frac{A}{\frac{1}{\tau} - g_L} (e^{-g_L t} - e^{-t/\tau})$$

and from Eq. (4), we have

$$m(t) = \left[(-g_L \mu(t) + \bar{I}^g + A e^{-t/\tau}) \frac{1}{P} e^{-(1-\mu(t))^2/2\sigma^2} \right]^+,$$

which, from Eq. (2b), results in a downstream synaptic current at $t = T$

$$\bar{I}_d^{\text{ff}} = S e^{-T/\tau} \int_0^T e^{t/\tau} m(t) \frac{dt}{\tau}. \quad (5)$$

After the end of the pulse, the current decays exponentially. This decaying current feeds forward and is integrated by the next layer. Thus, for an exact transfer, $\bar{I}_d^{\text{ff}}(S, \bar{I}^g, A, T) = A$.

In Fig. 4A, we show the bifurcation, found numerically in Fig. 3, for a range of values of S and \bar{I}^g . In Fig. 4B, we demonstrate the existence of a cusp catastrophe as a function of the width of the potential distribution σ . In Fig. 4C,D, we show how fast or slow (ghost) dynamics occur near the fold of the cusp depending on σ . In Fig. 4C,D, blue shades indicate locations where feedforward amplitudes are slowly changing, green and yellow shades where amplitudes change rapidly. Isoclines (the same in both panels) are plotted denoting the fixed point and nearby locations of small, but non-zero change. Note how in Fig. 4C, dynamics are fast and amplitudes rapidly approach the attractors, but in Fig. 4D, dynamics remain slow in a large region near the fold bifurcation. This region gives rise to an approximate line attractor.

In summary, to understand pulse-gated information propagation, we used a Fokker-Planck analysis to reveal that the apparent line attractor in the network of spiking neurons is associated with a cusp catastrophe as σ varies in state space. Locally, the nearly one-dimensional attracting manifold is




FIG. 4. σ -induced fold bifurcation ghost at cusp catastrophe. A) A plot of the surface of zeros of $\bar{I}^{\text{ff}}(S, \bar{I}^g, A, \sigma) - A$ plotted for (S, \bar{I}^g, A) with $\sigma = 1.0$. These are fixed points of the map of the synaptic feed forward current generated from Eqs. (5). Note the fold bifurcation across many values of \bar{I}^g , similar to that found numerically in Fig. 3. B) A plot of the surface of zeros of $\bar{I}^{\text{ff}}(S, \bar{I}^g, A, \sigma) - A$ plotted for (S, A, σ) with $\bar{I}^g = 6.5$. Note the obvious cusp catastrophe that occurs along the σ -axis. C) Zeros of $\bar{I}^{\text{ff}}(S, \bar{I}^g, A, \sigma) - A$ for $\bar{I}^g = 6.5$ and $\sigma = 0.2$ plotted on a colored background, where the color indicates the speed of departure toward stable fixed points (either upward above the unstable fixed point or downward below the unstable fixed point). For small σ , solutions rapidly approach the fixed points, giving a propagated binary code. D) Zeros of $\bar{I}^{\text{ff}}(S, \bar{I}^g, A, \sigma) - A$ for $\bar{I}^g = 6.5$ and $\sigma = 1.0$ again indicating the speed of departure toward stable fixed points. For σ near the cusp, there is a large region within which propagation away from the unstable point is slow, called a ghost. Ghosts occur near the cusp locus enabling graded propagation across many layers.

the result of a saddle-node bifurcation where the unstable fixed point is strongly attracting and weakly unstable. Dynamically, this reflects that the fold of this cusp separates a region of sigmoidal f-I curves from a region of bistability. Since this region occupies a sizeable volume in parameter space, the graded propagation is robust. Furthermore, the fact that a common cusp catastrophe underlies the fast timescale dynamics of our system indicates that this type of line attractor is generic and will persist in feedforward networks of other types of spiking neurons.

One of the major problems of neuroscience is to understand how complex neural functions emerge from the collective dynamics of neuronal networks. Towards this goal, researchers have tried to construct mechanical models using mean-field firing rate theories and large-scale numerical simulations. However, except for a few examples (e.g. [17–21]), the precise correspondence between the underlying microscopic spiking neurons and the macroscopic coherent dynamics of the neuronal populations has not been established.

Our analysis provides a significant step toward understanding how macroscopic attracting manifolds can emerge from the dynamic interactions of microscopic spiking neurons via the coordination of pulse-gating, synaptic weights, and

intrinsic and extrinsic noise (i.e. the distribution of the membrane potential across the population), and offers possible order parameters for which macroscopic descriptions can be derived from the underlying microscopic dynamical model. Furthermore, transfer mechanisms, such as those found in the graded and binary transfer parameter regimes shown in Fig. 4., provide a novel means of understanding dynamic network interactions such as the detailed measurements of population activity underlying complex neural tasks provided by modern experimental techniques. Already, these transfer mechanisms have been shown to be capable of implementing dynamic modules representing complex neural functions, such as short term memory, decision making, and the control of neural circuits [11, 12].

L.T. thanks the UC Davis Mathematics Department for its hospitality. A.T.S. would like to thank Liping Wei and the Center for Bioinformatics at the School of Life Sciences at Peking University for their hospitality. This work was supported by the Natural Science Foundation of China grants 91232715 (Z.X., L.T.), 91430216 (J.Z.), and U1530401 (J.Z.), by the Open Research Fund of the State Key Laboratory of Cognitive Neuroscience and Learning grant CNLZD1404 (Z.X., L.T.), by the Beijing Municipal Science and Technology Commission under contract Z151100000915070 (Z.X., L.T.), by the Undergraduate Honors Research Program of the School of Life Sciences at Peking University (Z.X., L.T.), and by the National Institutes of Health, CRCNS program NS090645 (A.T.S. and L.T.).

* Corresponding authors' email: ats@math.ucdavis.edu, taolt@mail.cbi.pku.edu.cn

- [1] X. J. Wang. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.*, 24(8):455–463, Aug 2001.
- [2] H. S. Seung. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. U.S.A.*, 93(23):13339–13344, Nov 1996.
- [3] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.*, 16(6):2112–2126, Mar 1996.
- [4] A. M. Bruno, W. N. Frost, and M. D. Humphries. Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. *Neuron*, 86(1):304–318, Apr 2015.
- [5] R. Ben-Yishai, R. L. Bar-Or, and H. Sompolinsky. Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 92(9):3844–3848, Apr 1995.
- [6] C. K. Machens, R. Romo, and C. D. Brody. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science*, 307(5712):1121–1124, Feb 2005.
- [7] V. Mante, D. Sussillo, K. V. Shenoy, and W. T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, Nov 2013.
- [8] P. E. Latham, S. Deneve, and A. Pouget. Optimal computation with attractor networks. *J. Physiol. Paris*, 97(4-6): 683–694, 2003.
- [9] S. Wu, D. Chen, M. Niranjana, and S. Amari. Sequential Bayesian decoding with a population of neurons. *Neural Comput*, 15(5):993–1012, May 2003.

- [10] S. Deneve, J. R. Duhamel, and A. Pouget. Optimal sensorimotor integration in recurrent cortical networks: a neural implementation of Kalman filters. *J. Neurosci.*, 27(21):5744–5756, May 2007.
- [11] A.T. Sornborger, Z. Wang, and L. Tao. A mechanism for graded, dynamically routable current propagation in pulse-gated synfire chains and implications for information coding. *J. Comput. Neurosci.*, August 2015. doi:10.1007/s10827-015-0570-8.
- [12] Z. Wang, A.T. Sornborger, and L. Tao. Graded, dynamically routable information processing with synfire-gated synfire chains. *PLoS Comp Biol*, 2016. doi: 10.1371/journal.pcbi.1004979.
- [13] T. Gisiger and M. Boukadoum. Mechanisms Gating the Flow of Information in the Cortex: What They Might Look Like and What Their Uses may be. *Front Comput Neurosci*, 5:1, 2011.
- [14] S. B. Floresco and A. A. Grace. Gating of hippocampal-evoked activity in prefrontal cortical neurons by inputs from the mediodorsal thalamus and ventral tegmental area. *J. Neurosci.*, 23(9):3930–3943, May 2003.
- [15] A. Hernandez, A. Zainos, and R. Romo. Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 97(11):6191–6196, May 2000.
- [16] C. Wang, Z.C. Xiao, Z. Wang, A.T. Sornborger, and L. Tao. A Fokker-Planck approach to graded information propagation in pulse-gated feedforward neuronal networks. *ArXiv*, (1512.00520), Dec 2015.
- [17] W. Gerstner. Time structure of the activity in neural network models. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics*, 51(1):738–758, Jan 1995.
- [18] D.Q. Nykamp and D. Tranchina. A population density approach that facilitates large-scale modeling of neural networks: Analysis and an application to orientation tuning. *J. Comput. Neurosci.*, 8:19–50, 2000.
- [19] D. Cai, L. Tao, A.V. Rangan, and D.W. McLaughlin. Kinetic theory for neuronal network dynamics. *Comm. Math. Sci.*, 4:97–127, 2006.
- [20] S. Ostojic and N. Brunel. From spiking neuron models to linear-nonlinear models. *PLoS Comput. Biol.*, 7(1):e1001056, Jan 2011.
- [21] E Montbrió, D. Pazo, and A. Roxin. Macroscopic description for networks of spiking neurons. *Phys. Rev. X*, 5:021028, 2015.